

Compressive Matched-Field Processing

William Mantzel, Justin Romberg, Karim Sabra *

August, 2011

Abstract

Source localization by matched-field processing (MFP) generally involves solving a number of computationally intensive partial differential equations. This paper introduces a technique that mitigates this computational workload by “compressing” these computations. Drawing on key concepts from the recently developed field of compressed sensing, it shows how a low-dimensional proxy for the Green’s function can be constructed by backpropagating a small set of random receiver vectors. Then, the source can be located by performing a number of “short” correlations between this proxy and the projection of the recorded acoustic data in the compressed space. Numerical experiments in a Pekeris ocean waveguide are presented which demonstrate that this compressed version of MFP is as effective as traditional MFP even when the compression is significant. The results are particularly promising in the broadband regime where using as few as two random backpropagations per frequency performs almost as well as the traditional broadband MFP, but with the added benefit of generic applicability. That is, the computationally intensive backpropagations may be computed offline independently from the received signals, and may be reused to locate any source within the search grid area.

1 Introduction

1.1 Background and Motivation

Matched field processing (MFP) continues to serve as one of the most widely used methods for localizing undersea targets acoustically. However, as the models governing undersea acoustic interactions become more sophisticated, often requiring fine-grain solutions to more complex partial differential equations, the tradeoff between run time and performance begins to worsen, perhaps unnecessarily. We will begin by discussing why this is the case and giving an overview of our approach to mitigate the problem.

MFP generalizes standard array beamforming methods (e.g. plane wave beamforming) for locating an acoustic source in a complex environment (such as a multipath shallow water waveguide). MFP has been studied extensively both theoretically and experimentally as described in several review articles [25, 2, 3]. MFP is usually implemented by systematically placing a test point source at each point of a spatial search grid of L candidate locations, computing the acoustic field (replicas) at all the elements of the receiver array and then correlating this modeled field with the data from the real point source whose localization

*[willem, jrom, ksabra3]@gatech.edu

is unknown to determine the best-fit location (see Fig. 1). This approach works well when that the computational replica environment is sufficiently accurate. However, this direct implementation of MFP using brute force search would require L computation runs which can become numerically cumbersome for large search space especially when simulating complex propagation environments.

One alternative to this direct implementation of MFP is to use a “backpropagation” algorithm (also referred to as “time-reversal imaging”) to locate the unknown source. In this case, a time-reversed version of the recorded data is used as an initial waveform excitation along the array aperture using the principle of superposition, and then subsequently “back-propagated” numerically in the replica environment towards the grid search area [17]. The unknown source location is then estimated from the maximum of the distribution of the backpropagated peak amplitude (or energy) across the grid search. Consequently, when compared to the direct implementation first mentioned, this backpropagation approach appears attractive at first glance, since it requires one computational run per unknown source. Nevertheless, this backpropagation approach becomes computationally expensive if multiple sources need to be located repetitively over the *same* search grid as the number of required computational runs would grow proportionally. For instance, this may occur when one tries to locate a source moving along a long track throughout the search space. Indeed, in order to be able locate any source throughout the search space using N receivers, MFP would require computing N backpropagations by using sequentially each individual receiver as a backpropagation source [25, 2]. This would allow determining the full set of Green’s functions associated with the channel between each search location and each receiver element. Alternatively, one could weight spatially the amplitude of the backpropagated signals along the receiver array using N different orthogonal codes (e.g. obtained from an Hadamard basis).

This article develops instead a compressive MFP formulation which reduces this computational burden by pre-computing the backpropagation of a number $M \ll N$ of random test signals. The results of these backpropagations effectively encode the Green’s function associated with the channel, and they can be re-used in subsequent localizations without any additional computational cost. This approach is inspired by recent work in the field of compressed sensing [6, 7, 15], whose central message is that random projections provide an effective encoding for sparse signals. The motivation for compressed sensing is typically concerned with reducing the cost of acquiring signals by shifting the workload from sensor hardware to software [8, 24, 20], and is natural in applications where physical measurements are expensive compared to numerical computations. Here we explore a variation on this theme: mitigating the computational workload in software instead of the sensing workload in hardware. The proposed compressive MFP allows us to estimate the underlying ambiguity function central to conventional MFP algorithms over the entire search space using only M computational runs instead of N , an effective speedup of a factor of N/M . In practice, these M simulations can be independently computed as a background process offline before the actual source signal is received.

1.2 Related Work

In this paper, we effectively demonstrate how classical localization procedures under a least-squares framework such as matched-field processing (MFP) may be solved in a reduced-dimensional space even without *a-priori* knowledge of the “best” dimension-reducing trans-

form. This property has been shown in similar forms in the mainstream canon of Compressed Sensing (CS) literature. Davenport et al. have described a number of useful variations on the theme of CS [14] including a matched filtering detector. They have also described the “smashed filter” that is designed primarily for classification between a finite number of sets, but could easily be extended to parametric estimation. Wakin has also established some rigorous results on parameter estimation that relate the recovery properties of a general compressive estimation problem to the properties of the manifold that these parameters induce. This work could be used to analyze this problem via its manifold parameters [26].

Carin et al. have utilized CS principles to show how a Green’s function of a scattering field that is compressible in the wavelet domain may be recovered from a small set of measurements, though they use incoherence in their structured measurements to recover a scattering field, while we primarily care about the location of the source [10]. Likewise, Marengo et al. have applied compressed measurements to the scattering problem, utilizing the target-sparse model to improve their performance [23].

Our work may also be viewed in the context of randomized SVDs [19]. In this field of research, the idea is to apply the matrix \mathbf{A} to a series of random vectors Φ_m as $\mathbf{A}\Phi_m$ in order to determine the range space of \mathbf{A} . For example, Chaillat et al. show how the inverse medium problem can be simplified using a dimension reducing random projection and solving the inverse problem in the reduced range-space [12]. Similarly to this field, we apply the time-reversal or adjoint of the Green’s function \mathbf{G}_w to random vectors in order to discover the range space of admissible ambiguity functions.

There is also a large amount of recent research performing multi-target tracking under the “target-sparse” assumption. That is, the methods propose to simultaneously localize several targets that lie on some grid (or generally some set of points) by solving an ℓ_1 minimization program. The recovered support resulting from this optimization corresponds to the grid points that the various targets are estimated to occupy. All of this work dovetails in very nicely with the main results of Compressed Sensing, which can be effectively leveraged to prove that the targets may be perfectly localized with high probability. Often, the painstaking effort in these papers involves showing that the Restricted Isometry Property (RIP) holds for the observation matrix. For example, Fannjiang et al. show the conditions under which a sufficiently small coherence is achieved for perfect recovery [16]. Gurbuz et al. show similar results for a *Compressive* beamformer, requiring a number of measurements on the order of the number of sources [18], but the application there is different in that they utilize a signal common to all sensors with an unknown time shift to localize their target in angle (assuming free space propagation), and apply the compression operator in time per-sensor instead of applying the operator across the range of sensors as we do. Also from a communications perspective, Cevher et al. demonstrate the relatively low amount of information to be transmitted for purposes of localization when using a Compressed Sensing framework [11]. These “target-sparse” approaches depend on targets lying exactly on the grid points. Also, by necessity these grid points must be spaced sufficiently far away from one another to avoid coherence-inducing correlations in the observation matrix. This creates a restrictive model of limited applicability. When a target is somewhere in between a set of grid points, the necessary conditions for recovery may not even *approximately* hold, similarly to how a discrete sinusoid corresponding to an off-grid point in the DFT will not be sparse in the frequency domain (or any other basis for any standard transform for that matter) due to DFT leakage. In contrast to this approach, we do not require our target




Figure 1: Schematic of a matched-field processing implementation in an ocean waveguide. The signal transmitted by a source (star symbol) located at an unknown location \vec{r}_0 is recorded along a N elements receiver array after multipath propagation. Using a computational model of the original ocean waveguide, the location \vec{r}_0 may be inferred by matching the actual received signals with the simulated replica waveforms obtained from varying the test source location (dot symbols) \vec{r} throughout the search grid area.

to lie on a grid point. However, instead of promising perfect recovery, we instead content ourselves to claim that our target may be localized to within a small neighborhood of the actual source location, or at least the location found via deterministic means.

1.3 Outline

The remainder of the paper is organized as follows. Section 2 briefly describes conventional MFP formulation for locating both single-frequency (narrowband) and broadband sources. Section 3 presents the corresponding compressive MFP (cMFP) formulation for both cases. Section 4 presents numerical simulations in a Pekeris waveguide [21, pg 540–552] illustrating the performance of cMFP in comparison to the conventional MFP results including the effects of additive ambient noise to the data and model mismatch due to uncertain knowledge of the actual environment. Section 5 extends this compressive approach to adaptive MFP. Section 6 summarizes the findings and conclusions drawn from this study.

2 Conventional MFP

A brief summary of the conventional MFP formulation is presented hereafter based on the standard solution of the linearized wave equation. The acoustic pressure field $y(\vec{r}, t)$ at a fixed point \vec{r} and time t produced by a point source located at \vec{r}_0 satisfies:

$$\frac{1}{c^2(\vec{r})} \frac{\partial^2 y(\vec{r}, t)}{\partial t^2} - \nabla^2 y(\vec{r}, t) = \alpha(t) \delta(\vec{r} - \vec{r}_0) \quad (1)$$

where $c(\vec{r})$ is the speed of sound and $\alpha(t)$ is the signal emitted by the source. The time-domain Green's function for the same environment $g(\vec{r}, \vec{r}_0, t)$ is, by definition, the solution of Eq. (1) for a impulsive point source (i.e. for $\alpha(t) = \delta(t)$) that satisfies all boundary

conditions [21, pg 540–552]. Using Eq. (1) (and assuming that the radiation condition applies as $\|\vec{r}\| \rightarrow \infty$) the Fourier transform of the recorded pressure field at \vec{r}_n , the n^{th} element of a receiver array ($n = 1..N$) (see Fig. 1), is denoted $y_\omega(\vec{r}_n)$ and given by:

$$y_\omega(\vec{r}_n) = \alpha_\omega g_\omega(\vec{r}_n, \vec{r}_0) \quad (2)$$

where ω is the frequency. The variables α_ω and $g_\omega(\vec{r}_n, \vec{r}_0)$ denote respectively the Fourier transform of the source signal and time-domain Green's function. Using vector notation, Eq. (2) can be restated as:

$$Y_\omega = \alpha_\omega G_\omega(\vec{r}_0), \quad (3)$$

where Y_ω is a $(N \times 1)$ column vector obtained by stacking the complex amplitudes $y_\omega(\vec{r}_n)$ measured along the receiver array. Similarly, the $(N \times 1)$ column vector $G_\omega(\vec{r})$ contains Green's functions $g_\omega(\vec{r}_n, \vec{r})$ between the N receiver array elements and a source located at \vec{r}_0 . Note that the position vectors are written in lowercase letters with arrows and the column vectors are written with capital letters in the remainder of this article.

2.1 Single-Frequency MFP

We start by considering the simplest MFP that works from measurements at a single frequency ω (as in (3)), known as the harmonic (or narrowband) formulation. Given a set of measurements $Y_\omega \in \mathbb{C}^N$ across the N receivers at frequency ω , we search for the location \vec{r} in our region of interest \mathcal{R} (and complex source amplitude β) that best accounts for these measurements by solving the least-squares problem

$$\arg \min_{\vec{r} \in \mathcal{R}} \min_{\beta \in \mathbb{C}} \|Y_\omega - \beta G_\omega(\vec{r})\|^2. \quad (4)$$

With the location \vec{r} fixed, the inner optimization problem is simply finding the closest point on the line spanned by $G_\omega(\vec{r})$ to the point Y_ω . Plugging in the closed-form solution to this problem (see Appendix A), the problem above reduces to:

$$\arg \min_{\vec{r}} \|Y_\omega\|^2 - \frac{|Y_\omega^H G_\omega(\vec{r})|^2}{\|G_\omega(\vec{r})\|^2} = \arg \max_{\vec{r}} \frac{|Y_\omega^H G_\omega(\vec{r})|^2}{\|G_\omega(\vec{r})\|^2}, \quad (5)$$

(where Y_ω^H denotes the Hermitian transpose) which we will refer to as the *normalized ambiguity function*, and will refer to its maximization as *normalized Matched Field Processing* (nMFP). We show an example of the normalized ambiguity function in Fig. 2.a.

The term $Y_\omega^H G_\omega(\vec{r})$ can be computed at every location \vec{r} in an efficient manner using time-reversal. Precise values for $\|G_\omega(\vec{r})\|^2$ are typically not available when computing the backpropagation $Y_\omega^H G_\omega(\vec{r})$. However it is often the case (and we will assume this here) that these energies either do not vary much across our locations of interest, or vary predictably (e.g. cylindrical spreading of the field amplitude). Dropping the denominator yields the so-called *unnormalized ambiguity function* (alternatively the unnormalized Bartlett formulation) [25, 2, 3], the objective function used for estimating the source location:

$$\hat{\vec{r}} = \arg \max_{\vec{r}} |h(\vec{r})|^2 \quad \text{where} \quad h(\vec{r}) = Y_\omega^H G_\omega(\vec{r}), \quad (6)$$

2.2 Broadband MFP

Now suppose that most of the energy of the source signal occupies some continuous bandwidth $[\omega_{\min}, \omega_{\max}]$, known as the broadband formulation. Ideally, we would solve (4) over a continuum of ω values. However, for the sake of source localization, it is computationally advantageous to sample this bandwidth at K frequencies $\omega_1, \omega_2, \dots, \omega_K$, yielding K measurement vectors Y_{ω_k} where $k \in \{1, 2, \dots, K\}$. In this way, we can achieve a computational complexity at most K times the single frequency case, without sacrificing much precision.

We now search for the location \vec{r} that jointly matches the joint behavior of the measurements Y_{ω} over multiple frequencies $\omega_1, \omega_2, \dots, \omega_K$. The least-squares problem from (4) becomes

$$\arg \min_{\vec{r}} \min_{\beta_{\omega_1}, \dots, \beta_{\omega_K}} \sum_{k=1}^K \|Y_{\omega_k} - \beta_{\omega_k} G_{\omega_k}(\vec{r})\|^2. \quad (7)$$

The inner optimization problem is separable over the β_{ω_k} , and so the above is equivalent to

$$\arg \min_{\vec{r}} \sum_{k=1}^K \min_{\beta_{\omega_k}} \|Y_{\omega_k} - \beta_{\omega_k} G_{\omega_k}(\vec{r})\|^2 = \arg \max_{\vec{r}} \sum_{k=1}^K \frac{|Y_{\omega_k}^H G_{\omega_k}(\vec{r})|^2}{\|G_{\omega_k}(\vec{r})\|^2}. \quad (8)$$

As before, if the energies $\|G_{\omega_k}(\vec{r})\|^2$ are homogenous across space and frequency, then a reasonable unnormalized approximation to the above is

$$\arg \max_{\vec{r}} \sum_{k=1}^K |Y_{\omega_k}^H G_{\omega_k}(\vec{r})|^2 = \sum_{k=1}^K |h_{\omega_k}(\vec{r})|^2. \quad (9)$$

The formulation in (9) assumes that the source amplitudes β_{ω_k} are unknown. If we have knowledge of the source signal's complex amplitudes, that is we know them up to a common amplitude and phase, then (7) can be refined to

$$\arg \min_{\vec{r}} \min_{\beta \in \mathbb{C}} \left\| \begin{bmatrix} Y_{\omega_1} \\ Y_{\omega_2} \\ \vdots \\ Y_{\omega_K} \end{bmatrix} - \beta \begin{bmatrix} \alpha_{\omega_1} G_{\omega_1}(\vec{r}) \\ \alpha_{\omega_2} G_{\omega_2}(\vec{r}) \\ \vdots \\ \alpha_{\omega_K} G_{\omega_K}(\vec{r}) \end{bmatrix} \right\|^2 \quad (10)$$

where the source amplitudes α_{ω_k} are fixed and known. Applying again the results from Appendix A, the inner optimization program can be solved in closed form, and so (10) is equivalent to

$$\arg \max_{\vec{r}} \frac{\left| \sum_{k=1}^K \alpha_{\omega_k} Y_{\omega_k}^H G_{\omega_k}(\vec{r}) \right|^2}{\sum_{k=1}^K |\alpha_{\omega_k}|^2 \|G_{\omega_k}(\vec{r})\|^2}, \quad (11)$$

as shown in Fig. 2.b, which we can approximate (by removing the denominator) as its unnormalized counterpart

$$\arg \max_{\vec{r}} \left| \sum_{k=1}^K \alpha_{\omega_k} h_{\omega_k}(\vec{r}) \right|^2. \quad (12)$$

Hereafter, we will refer to (8) and (9) as the *incoherent* MFP formulation, and (11) and (12) as the *coherent* MFP formulation.




Figure 2: Single-frequency (left column) and broadband-coherent (right column) ambiguity functions. These ambiguity functions shown on the dB scale ($20 \log_{10}(\cdot)$) for: (a, b) the standard MFP as described in Eqs. (5) and (11), and (c, d) cMFP as described in Eqs. (15) and (22) for the single-frequency case with $M = 10$ and broadband coherent case with $M = 2$ measurements per frequency, and (e, f) cMFP for the single-frequency case with $M = 30$ and broadband coherent case with $M = 20$ measurements per frequency.

3 Compressive MFP

In this section, we describe Compressive Matched Field Processing (cMFP). This is an efficient method for acquiring a *compressed* version of the Green's function operator $G_\omega(\vec{r})$ that exhibits a behavior in some regards similar to the dimension-reduced counterpart achieved via Principal Component Analysis, but may be obtained with only incomplete knowledge of the Green's function $G_\omega(\vec{r})$. Our approach works by precomputing the backpropagation of a small number of hypothetical received signals to construct a dimension-reduced proxy for the Green's function. Then, given the actual observed data Y_ω we localize the source by finding the closest match between the received signal and the Green's function in the compressed domain. With the compressed version of $G_\omega(\vec{r})$ in hand, locating the source only requires computing a series of short inner products. In addition, the compressed version of $G_\omega(\vec{r})$ is independent of the received signal, and so can be pre-computed and re-used for later observations. As we will demonstrate in Section 4, this cMFP strategy is effective even when the number of pre-computed random backpropagations is far fewer than what would be required for a full acquisition of $G_\omega(\vec{r})$ over the whole search grid area.

3.1 Single-Frequency cMFP

We start by discussing the single-frequency case in detail. First, we compute the *compressed* Green's function $\mathbf{\Phi}G_\omega(\vec{r})$, where $\mathbf{\Phi}$ is a $M \times N$ *encoding matrix*. Note that matrices are written in boldface letters in the remainder of this article. We construct $\mathbf{\Phi}G_\omega(\vec{r})$ by backpropagating (i.e. applying G_ω^H to) a series of *test vectors* $\Phi_1, \dots, \Phi_M \in \mathbb{C}^N$ — we will discuss how the Φ_m are chosen in the next section.

The result of one of these computations $\Phi_m^H G_\omega(\vec{r})$ is a complex-valued acoustic field over \vec{r} and requires as much effort to compute as the ambiguity function $h(\vec{r})$. We stack up the results of these precomputations as rows in the ensemble

$$\begin{bmatrix} \Phi_1^H G_\omega(\vec{r}) \\ \Phi_2^H G_\omega(\vec{r}) \\ \vdots \\ \Phi_M^H G_\omega(\vec{r}) \end{bmatrix} = [\Phi_1 \quad \Phi_2 \quad \dots \quad \Phi_M]^H G_\omega(\vec{r}) = \mathbf{\Phi}G_\omega(\vec{r}). \quad (13)$$

This ensemble gives us access to an indirect, dimension-reduced version of $G_\omega(\vec{r})$.

Given observations Y_ω , we search for the \vec{r} that best explains these random backpropagations in the compressed space. The least-squares program (4) becomes

$$\arg \min_{\vec{r}} \min_{\beta} \|\mathbf{\Phi}Y_\omega - \beta \mathbf{\Phi}G_\omega(\vec{r})\|^2, \quad (14)$$

which, again using the results from Appendix A, reduces to

$$(\text{narrowband cMFP}) \quad \arg \max_{\vec{r}} \frac{|Y_\omega^H \mathbf{\Phi}^H \mathbf{\Phi}G_\omega(\vec{r})|^2}{\|\mathbf{\Phi}G_\omega(\vec{r})\|^2} \quad \text{where} \quad \tilde{h}(\vec{r}) = Y_\omega^H \mathbf{\Phi}^H \mathbf{\Phi}G_\omega(\vec{r}). \quad (15)$$

The function $\tilde{h}(\vec{r})$ is shown in Fig. 2.c and 2.e, and can be interpreted as a compressed version of the ambiguity function $h(\vec{r})$ in (6) shown in Fig. 2.a. The cross sections in range and depth of these ambiguity functions are shown in Fig. 3.a and 3.b.

Note that unlike the standard MFP, in this case the precomputations give us direct access to the denominator $\|\Phi G_\omega(\vec{r})\|^2$ (we simply take the norms of the columns of $\Phi G_\omega(\vec{r})$), and so we leave it in the optimization program. As shown in the results section, this normalization term plays an important role in improving the source location estimation when the magnitude of the Green's function varies significantly across the search grid area.

Notice that an evaluation of (15) at a point \vec{r} essentially only requires an inner product between the encoded observations ΦY_ω and the M -vector $\Phi G_\omega(\vec{r})$ formed from the back-propagated fields at point \vec{r} from all M test vectors.

3.2 Random Projections

The question remains as to how to choose the encoding matrix Φ so that solution to the cMFP (15) is the same (or close to) the solution to the standard MFP (6). The corresponding least-squares problems are

$$\text{standard MFP :} \quad \arg \min_{\vec{r}, \beta} \|Y_\omega - \beta G_\omega(\vec{r})\|^2 \quad (16)$$

$$\text{cMFP :} \quad \arg \min_{\vec{r}, \beta} \|\Phi (Y_\omega - \beta G_\omega(\vec{r}))\|^2. \quad (17)$$

These two programs will have similar solutions if their functionals are close to one another for all values of β and \vec{r} . If $Y_\omega = \alpha G_\omega(\vec{r}_0)$, then the performance of the cMFP will match that of the standard MFP when Φ preserves the energy of the differences between the observations Y_ω and all scalar multiples of the Green's function at different points:

$$\|\Phi(F_1 - F_2)\|^2 \approx \|F_1 - F_2\|^2 \quad \text{for all } F_1, F_2 \in \mathcal{F} := \{F : F = \alpha G_\omega(\vec{r}), \alpha \in \mathbb{C}; \vec{r} \in \mathcal{R}\}. \quad (18)$$

Essentially, we want Φ to *stably embed* (i.e. preserve the distances between members of) the set \mathcal{F} into \mathbb{C}^M .


We propose taking Φ to be a random linear mapping. This choice is inspired both by classical results in theoretical computer science and from the recently developed theory of compressive sensing. In the mid-1980s, Johnson and Lindenstrauss [22] demonstrated that the distances within a finite set of n points are essentially preserved through a random projection into a space of dimension $\sim \log n$ (see also [13, 1]). Recently it has been shown that this same type of projection also embeds sparse signals into a low-dimensional subspace [4], a result which plays a key role in compressive sampling [7, 9], and are effective at reducing the dimensionality of certain types of manifolds [5].

We will discuss the particular the case where Φ is a *random orthoprojection*, although the results will be almost identical for many different choices of random Φ (e.g. with entries that are independent and identically distributed Gaussian or ± 1 random variables). To generate Φ , we simply draw an $M \times N$ matrix of independent Gaussian random variables with unit variance, orthonormalize the rows using the Gram-Schmidt (or QR) algorithm, and then multiply by $\sqrt{N/M}$. For an arbitrary fixed vector F , the random orthoprojection Φ obeys two properties[13]:


$$\mathbb{E} [\|\Phi F\|^2] = \|F\|^2 \quad (19)$$

$$\mathbb{P} \{ |\|\Phi F\|^2 - \|F\|^2| > \epsilon \} \leq 2e^{-\frac{M}{2\|F\|^2}(\epsilon^2/2 - \epsilon^3/3)}. \quad (20)$$


This allows us to interpret the compressed energy functional $\|\Phi(Y_\omega - \beta G_\omega(\vec{r}))\|^2$ in (17) as a random process, indexed by β and \vec{r} , whose mean is the standard energy functional $\|Y_\omega -$




(a)



(b)



(c)



(d)

Figure 3: Cross-sections of the ambiguity functions displayed on Fig. 2: (a, b) single-frequency case described by Eqs. (5) and Eqs. (15); (c, d) broadband coherent case Eqs. (11) and (22); range (left column) and depth (right column). Here we show the normalized standard MFP (nMFP) and the cMFP (cMFP) for various values of M . The dashed lines show the boundaries for the main lobe and region of uncertainty that we are able to localize within under the presence of modest noise.

$\beta G_\omega(\vec{r})\|^2$ in (16). At a fixed point β, \vec{r} , this random process is concentrated around its mean roughly like a Gaussian random variable with standard deviation $\sqrt{2/M}\|Y_\omega - \beta G_\omega(\vec{r})\|$. The larger we make M (the more random vectors we precompute backpropagations for), the tighter the concentration. By construction, when $M = N$, $\Phi^H \Phi = I$ and we have acquired a “lossless” version of the Green’s function $G_\omega(\vec{r})$, meaning that the functionals are exactly equal to one another. In general, however, we will be interested in cases where there is a significant compression factor $M \ll N$ and benefit from the associated computational savings.

3.3 Broadband cMFP

The cMFP formulation can be readily extended to combine observations at multiple frequencies in both the incoherent and coherent cases. For frequencies $\omega_1, \omega_2, \dots, \omega_K$, we generate a sequence of $M \times N$ random matrices $\Phi_{\omega_1}, \Phi_{\omega_2}, \dots, \Phi_{\omega_K}$ and backpropagate the rows of each (for a total of MK time-reversals) to acquire $\Phi_{\omega_1} G_{\omega_1}(\vec{r}), \dots, \Phi_{\omega_K} G_{\omega_K}(\vec{r})$. Then given observations $Y_{\omega_1}, \dots, Y_{\omega_K}$, we compress them by calculating $\Phi_{\omega_1} Y_{\omega_1}, \dots, \Phi_{\omega_K} Y_{\omega_K}$, and then using the compressed versions of the G_{ω_k} , we proceed as in (7) for the incoherent case

$$\begin{aligned}
(\text{incoherent cMFP}) \quad & \arg \min_{\vec{r}} \min_{\beta_{\omega_1}, \dots, \beta_{\omega_K}} \sum_{k=1}^K \|\Phi_{\omega_k} Y_{\omega_k} - \beta_{\omega_k} \Phi_{\omega_k} G_{\omega_k}(\vec{r})\|^2 \\
& = \arg \min_{\vec{r}} \sum_{k=1}^K \min_{\beta_{\omega_k}} \|\Phi_{\omega_k} Y_{\omega_k} - \beta_{\omega_k} \Phi_{\omega_k} G_{\omega_k}(\vec{r})\|^2 \\
& = \arg \max_{\vec{r}} \sum_{k=1}^K \frac{|Y_{\omega_k}^H \Phi_{\omega_k}^H \Phi_{\omega_k} G_{\omega_k}(\vec{r})|^2}{\|\Phi_{\omega_k} G_{\omega_k}(\vec{r})\|^2}, \\
& = \arg \max_{\vec{r}} \sum_{k=1}^K \frac{|\tilde{h}_{\omega_k}(\vec{r})|^2}{\|\Phi_{\omega_k} G_{\omega_k}(\vec{r})\|^2} \tag{21}
\end{aligned}$$

and as in (10) for the coherent case:

$$\begin{aligned}
(\text{coherent cMFP}) \quad & \arg \min_{\vec{r}} \min_{\beta} \left\| \begin{bmatrix} \Phi_{\omega_1} Y_{\omega_1} \\ \Phi_{\omega_2} Y_{\omega_2} \\ \vdots \\ \Phi_{\omega_K} Y_{\omega_K} \end{bmatrix} - \beta \begin{bmatrix} \alpha_{\omega_1} \Phi_{\omega_1} G_{\omega_1}(\vec{r}) \\ \alpha_{\omega_2} \Phi_{\omega_2} G_{\omega_2}(\vec{r}) \\ \vdots \\ \alpha_{\omega_K} \Phi_{\omega_K} G_{\omega_K}(\vec{r}) \end{bmatrix} \right\|^2 \\
& = \arg \max_{\vec{r}} \frac{\left| \sum_{k=1}^K \alpha_{\omega_k} Y_{\omega_k}^H \Phi_{\omega_k}^H \Phi_{\omega_k} G_{\omega_k}(\vec{r}) \right|^2}{\sum_{k=1}^K |\alpha_{\omega_k}|^2 \|\Phi_{\omega_k} G_{\omega_k}(\vec{r})\|^2} \\
& = \arg \max_{\vec{r}} \frac{\left| \sum_{k=1}^K \alpha_{\omega_k} \tilde{h}_{\omega_k}(\vec{r}) \right|^2}{\sum_{k=1}^K |\alpha_{\omega_k}|^2 \|\Phi_{\omega_k} G_{\omega_k}(\vec{r})\|^2}. \tag{22}
\end{aligned}$$

The incoherent and coherent case are respectively illustrated in Fig. 2.d and 2.f and in Fig. 3.c and 3.d. Note that in this coherent case, the optimization is identical in its structure

to the single-frequency case. In particular, by concatenating:

$$G(\vec{r}) = \begin{bmatrix} \alpha_{\omega_1} G_{\omega_1}(\vec{r}) \\ \alpha_{\omega_2} G_{\omega_2}(\vec{r}) \\ \vdots \\ \alpha_{\omega_K} G_{\omega_K}(\vec{r}) \end{bmatrix} \quad Y = \begin{bmatrix} Y_{\omega_1} \\ Y_{\omega_2} \\ \vdots \\ Y_{\omega_K} \end{bmatrix} \quad \Phi = \begin{bmatrix} \Phi_{\omega_1} & & & \\ & \Phi_{\omega_2} & & \\ & & \ddots & \\ & & & \Phi_{\omega_K} \end{bmatrix}, \quad (23)$$

we see that the coherent broadband formulation (22) shares the same formulation as the single frequency case (15).

4 Numerical simulations

In this section, we present numerical experiments demonstrating that underwater acoustic sources can be localized from highly compressed versions of the Green's functions. Our cMFP results give locations estimates for single-frequency, incoherent broadband, and coherent broadband that are comparable with the traditional MFP. After the initial pre-computation (which consists of backpropagating the random codes at each frequency), the cMFP is substantially faster than the traditional MFP, requiring only a short inner product to be calculated at each search location.

The MATLAB code generating all the numerical results presented in this section is available online ¹.

4.1 Numerical set-up

All numerical simulations were conducted using a 200m deep Pekeris waveguide and the Green's functions were computed using a standard normal mode code [21, pg 540–552]. The two dimensional search grid area in depth and range spans respectively [10m 190m], and [5000m 5810m] for the single frequency and broadband incoherent simulations. The range span for the broadband coherent simulations was reduced to [5000m 5270m] to keep constant the number of search locations over which the ambiguity functions are computed since the effective resolution of the ambiguity function in the coherent case was about 3 times higher in range (see Fig. 2 and Fig. 3). A uniformly spaced vertical line array with $N = 37$ elements spaced between 10 and 190 meters was used to sample the acoustic field. The Green's functions between each of the search locations and the receiver array (see Fig. 1) were calculated across $K = 20$ different frequencies between 141 Hz and 160 Hz (the narrowband configuration uses 150 Hz). Given the selected numerical set-up, the natural resolution in frequency of the computed Green's function is around 5 Hz; that is, $G_{\omega_1}(\vec{r})$ and $G_{\omega_2}(\vec{r})$ are essentially uncorrelated when $|\omega_1 - \omega_2| \geq 10\pi$. The selected sample spacing of 1 Hz falls well within this frequency resolution.

After selected a source location \vec{r}_0 inside the region of interest, observations at the K frequencies were simulated using the forward model, and uncorrelated zero-mean Gaussian noise was added to the result:

$$Y_{\omega_k} = \alpha_{\omega_k} G_{\omega_k}(\vec{r}_0) + Z_k, \quad Z_k \in \mathbb{C}^N, \quad Z_k \sim \text{Normal}(0, \sigma^2 I), \quad (24)$$

¹Download the code at <http://users.ece.gatech.edu/~wmantzel3/cmfp/code.zip>.




Figure 4: Definition of the elliptical distance metric used for the performance study of cMFP. Because range errors tend to be greater than depth errors in long-range localization estimates, our results use an elongated distance metric that gives greater weight to the depth than the range, leading to unit balls that are ellipses instead of circles as shown here (see Eq. (27)). The color scheme for this ambiguity surface $10 \log_{10}(|h(\vec{r})|^2) = 10 \log_{10}(|Y_\omega^H G_\omega(\vec{r})|^2)$ has been lightened somewhat to allow for better visibility for the overlaid ellipses.

where each Z_k has i.i.d. Gaussian real and imaginary parts with variance $\sigma^2/2$. In all of our experiments, we set $\alpha_{\omega_k} = 1$ for all k . The signal-to-noise ratio (SNR) corresponding to noise variance σ^2 is

$$\text{SNR} = 10 \log_{10} \left(\frac{|\alpha_\omega|^2 \|G_\omega(\vec{r}_0)\|^2}{N\sigma^2} \right) \quad (25)$$

in the single frequency case, and

$$\text{SNR} = 10 \log_{10} \left(\frac{\sum_{k=1}^K |\alpha_{\omega_k}|^2 \|G_{\omega_k}(\vec{r}_0)\|^2}{KN\sigma^2} \right) \quad (26)$$

in the broadband case. Unless otherwise stated, we used an SNR of 16 dB.

Given a set of observations, we estimate the source location by solving (15) (single frequency), (21) (broadband incoherent), or (22) (broadband coherent) and compare against the standard MFP formulations (6), (9), and (12). As stated, these optimizations problems are over a continuous variable \vec{r} — in practice, we compute these functionals on a finite grid of points and choose the maximum from amongst these points. We used a 90×90 grid for the simulations presented below, which corresponds to 2m spacing in depth, a 9m spacing in range in the single-frequency and broadband incoherent cases, and 3m spacing in range in the broadband coherent case. We wish to emphasize that while our solution will of course lie on one of these grid points, the actual source location is chosen to be an arbitrary point.

The natural resolutions in depth and range of the ambiguity function $h_\omega(\vec{r})$ differ, as shown in Fig. 4 and Fig. 3 for a source located at $\vec{r}_0 = (5540\text{m}, 100\text{m})$ in (range, depth) for

a single frequency $\omega = 300\pi$ rad/sec (150 Hz) for a source located at $\vec{r}_0 = (5540, 100)$ in (range, depth). In this case, the main lobe has a width of ~ 360 m in range and ~ 32 m in depth. Again the grid spacing of $9m/2m$ in range/depth falls well within this resolution. The spatial resolution of the ambiguity surface in the selected multi-modal Pekeris waveguide is primarily a function of the source-receiver array configuration as well as the selected frequency band [25, 2]. In light of these differing spatial resolutions, we use a weighted norm to report distance errors in most cases presented here in this section. The distance from a point $\vec{r}_0 = (r_0^{\text{range}}, r_0^{\text{depth}})$ to the estimated source location $\hat{\vec{r}}$ is computed using the elliptical distance:

$$\|\vec{r}_0 - \hat{\vec{r}}\|_e = \sqrt{\left(\frac{r_0^{\text{range}} - \hat{r}^{\text{range}}}{e^{\text{range}}}\right)^2 + \left(\frac{r_0^{\text{depth}} - \hat{r}^{\text{depth}}}{e^{\text{depth}}}\right)^2}. \quad (27)$$

We use $e^{\text{depth}} = 3\text{m}$ and $e^{\text{range}} = 36\text{m}$ for the single-frequency and incoherent cases, and $e^{\text{range}} = 12\text{m}$ in the coherent case. The values of e^{depth} and e^{range} were chosen so that the contour $\{\vec{r} : \|\vec{r}_0 - \vec{r}\|_e = 1\}$ was approximately the same as the isosurface of the ambiguity function at 0.9 of its maximum. Equidistant points from $\vec{r}_0 = (5540, 100)$ for $\|\vec{r}_0 - \vec{r}\|_e = 1, 5$, and 10 are shown in Fig. 4. For example, an error of 14.4 meters in range and 0.9 meters in depth translates to 0.5 units of distance error in the elliptical $\|\cdot\|_e$ norm.

4.2 Localization performance of cMFP.

Fig. 5a compares the performance of cMFP (see Eq. (15)) and MFP (see Eq. (eq:amb-norm-eq:amb)) for locating a harmonic source ($f = 150\text{Hz}$). The SNR of the received data vector (see Eq. 24) was set to 16 dB. For a fixed M we aggregate performance statistics across 1000 simulations: 100 different source locations (chosen from \mathcal{R} uniformly at random) and 10 different draws of the Φ_ω for each location. For each test simulation, the error between the true and estimated target location was recorded in units of the target ellipse radius (see Fig. 4). From the results of the 1000 simulations, we calculated the empirical distance tail probability $P_M(d)$ -for a given number of random backpropagations M - as the fraction of results that produces a location estimate $\hat{\vec{r}}$ with $\|\hat{\vec{r}} - \vec{r}_0\|_e > d$. As shown, we are able to estimate the target within the unit ellipse more than 99% of the time from only $M = 6$ test vectors. Notice that the cMFP actually outperforms the unnormalized version of the MFP (from (6) above) when $M \approx 6$. This happens because the cMFP has an estimate of the normalizing factor in the denominator, as shown in (15). The cMFP is really an estimate of the normalized MFP in (5), and indeed that formulation is what the cMFP approaches as the number of random backpropagations M becomes equal to number of receivers N .

The cMFP was also tested in a variety of SNR for the single-frequency case. Fig. 5b shows the probability that the localization estimate is within the first ellipse (i.e. $d < 1$) as a function of the number of random backpropagations M . In all cases, the failure probability asymptotically decreases exponentially in the number of random backpropagations. Finally, Fig. 5c shows the tail probability of distance error for a fixed number of random backpropagations $M = 20$. As expected, the performance of cMFP gradually decrease as the SNR of the measurements is reduced from 16dB to 0dB, similarly to what occurs when using conventional MFP [3].

Fig. 6 and Fig. 7 show a similar performance study for respectively the broadband incoherent









Figure 5: (a) Tail probability of distance error $\|\hat{\vec{r}} - \vec{r}_0\|_e$ (see Eq. (27)) for the single-frequency cMFP formulation (see Eq. (15)) at 150 Hz. $P_M(d)$ is the probability that the localization is worse than some distance d using M random backpropagations. The dashed lines indicate the performance under normalized and unnormalized MFP (Eq. (5) and Eq. (6)). The next two plots show results for $P_M(d)$ over various SNRs of the received signal with (b) fixing $d = 1$ and (c) fixing $M = 20$.



(a)




(b)




(c)


Figure 6: Same as Fig. 5 but using instead the incoherent broadband cMFP formulation (see Eq. (22))



(a)



(b)



(c)

Figure 7: Same as Fig. 5 but using instead the coherent broadband cMFP formulation (see Eq. (22))

cMFP (see Eq. (21)) or broadband coherent cMFP (see Eq. (22)) formulations, including the influence of the SNR of the measurements as well as the number of random backpropagations M . Note that in Fig. 7, the horizontal axis is normalized differently than in the other two cases due to the different spatial resolutions of the ambiguity surfaces (see Eq. (27)). Our intentions are not to directly compare the performance of each of the three cMFP formulations (the coherent localization being always better as expected), but rather to show that in each case the selected cMFP formulation performs as well as the corresponding normalized MFP formulation and better the corresponding unnormalized MFP formulation. This is especially true for the broadband coherent cMFP results as Fig. 7.a shows that with just $M = 1$ measurement per frequency, we achieve an error within 3 times what standard MFP gives us at least 90% of the time, and with $M = 2$, we fall within about 10% distance error of what MFP gives us about 99% of the time.

Furthermore, note that we do not show results for $M = 1$ for the broadband incoherent cMFP formulation (Fig. 6.a) as in this case $\Phi_k G_{\omega_k}(\vec{r})$ is a scalar for each \vec{r} , and (21) reduces to

$$\arg \max_{\vec{r}} \sum_{k=1}^K \frac{|Y_{\omega_k}^H \Phi_k^H \Phi_k G_{\omega_k}(\vec{r})|^2}{\|\Phi_k G_{\omega_k}(\vec{r})\|^2} = \arg \max_{\vec{r}} \sum_{k=1}^K \frac{|\Phi_k Y_{\omega_k}|^2 |\Phi_k G_{\omega_k}(\vec{r})|^2}{|\Phi_k G_{\omega_k}(\vec{r})|^2} \quad (28)$$

$$= \arg \max_{\vec{r}} \sum_{k=1}^K |\Phi_k Y_{\omega_k}|^2. \quad (29)$$

This optimization problem is ill-defined, as the functional does not depend on \vec{r} .

4.3 Evolution of the main lobe to side lobe ratio of the cMFP ambiguity surface.

Fig. 8 shows the logarithmic variations of the main lobe to side lobe ratio of the ambiguity surface obtained with the single frequency and broadband coherent cMFP formulations for increasing number of random backpropagations M . In each case, the displayed values represent the median value of the main lobe to side lobe ratios obtained from 1000 simulations for each value of M . Here the main lobe is defined as the maximum of the ambiguity surface $|h(\vec{r})|$ (obtained from the corresponding conventional MFP formulation, e.g. see Fig. 2a-b and Fig. 3) over the region of interest \mathcal{R} , and the side lobe as the maximum of $|h(\vec{r})|$ over the search area \mathcal{R} excluding an ellipse E of the approximate size of the main lobe. We show the cross sections of the ambiguity function in Fig. 3 where we illustrate our choice of main lobe ellipse parameters that define our main lobe ellipse E . For the single frequency case, the ellipse has parameters $e^{\text{range}} = 180$ meters and $e^{\text{depth}} = 16$ meters (the broadband coherent case uses $e^{\text{range}} = 72$ meters and $e^{\text{depth}} = 16$ meters) as illustrated in Fig. 3. The logarithmic value of the main lobe to side lobe ratio is computed as:

$$20 \log_{10} \left(\frac{\max_{\vec{r} \in \mathcal{R}} |h(\vec{r})|}{\max_{\vec{r} \in \mathcal{R} \setminus E} |h(\vec{r})|} \right), \quad (30)$$

Note that for small M values, the cMFP side lobes may be significantly larger than their standard MFP counterparts. The concentration inequality (20) suggests that as M gets larger, the side lobes dampen. This behavior is observed in Fig. 8. Note that since the Φ

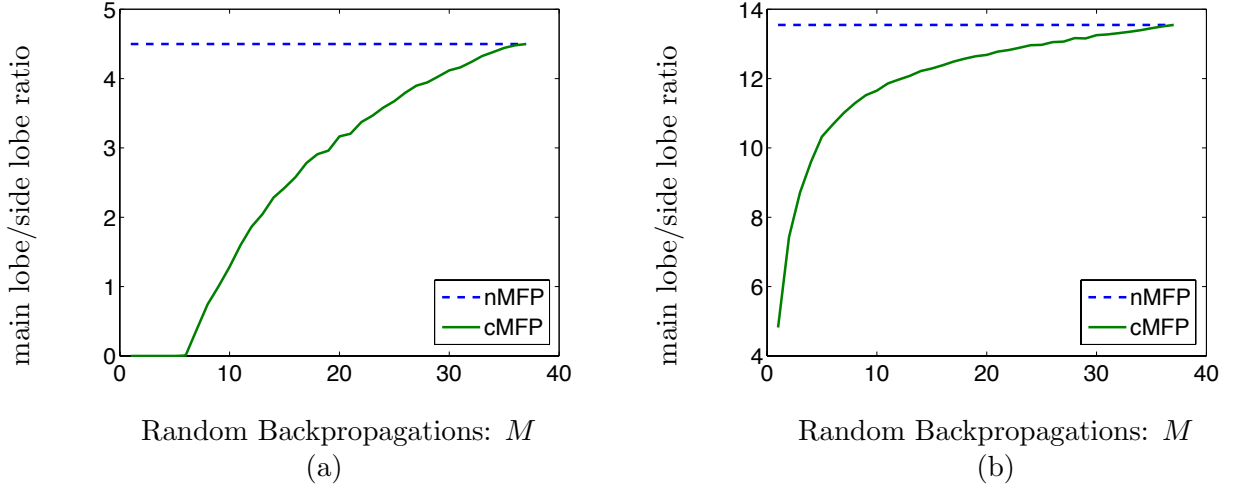


Figure 8: Evolution of the main lobe to side lobe ratio (in dB) of the estimated ambiguity surface (e.g. see Fig. 2) vs. number of random backpropagations M using either (a) the single frequency cMFP formulation at 150 Hz or (b) the broadband coherent MFP formulation (see Eq. (22)). Note that in each case the main lobe to side lobe ratio of the ambiguity surface obtained with cMFP reaches the main lobe to side lobe ratio value obtained using the corresponding nMFP formulation (dashed line) when $M = N = 37$.

matrix is an isometry when $M = N$, the side lobes in this case are exactly the same as for the standard MFP.

4.4 Influence of model mismatch on the cMFP performance

Previous studies have shown extensively that one major liability of MFP is sensitivity to model mismatch which occurs when one has an incorrect model for the ocean waveguide (e.g. sound speed profile error) [2]. Since MFP exploits the knowledge of the environment (via the Green's functions), its numerical accuracy must be sufficiently accurate, to ensure accurate source localization. Here we simply ensure that the localization accuracy of cMFP remains comparable to conventional MFP in the presence of error in the sound speed value. To do so, a set of received signals with a set SNR of 16 dB were computed for a reference sound speed of 1520 m/s. The broadband coherent cMFP -using $M = 4$ random backpropagations per frequency (see (22)) and normalized MFP formulation (see (11)) were then implemented using backpropagations in a simulated environment with different nominal values for the sounds speed (between 1520 m/s and 1530 m/s) than the reference value of 1520 m/s. Fig. 9 shows that the cMFP performs substantially the same as traditional MFP, for better or for worse. We show the average distance error in actual Euclidean distance (meters) as

$$\sqrt{(r_0^{\text{range}} - r^{\text{range}})^2 + (r_0^{\text{depth}} - r^{\text{depth}})^2}, \quad (31)$$

instead of ellipse distance. The small localization error occurring even without modeling error is due to the fact that the true source location did not coincide exactly with one the grid search location \vec{r} .

Note also that the range error tends to dominate for sufficiently large modeling error: the




Figure 9: Evolution of the localization error for broadband coherent cMFP and corresponding conventional MFP a for increasing error of the modeled sound speed value. The correct sound speed value is 1520m/s here. Notice that the localization errors obtained from cMFP (circle symbols) match closely the localization errors obtained from standard MFP (cross symbols).

slope of the displayed error values is roughly $5000/1520$ (the nominal range divided by the nominal speed of sound) as we would expect because a 15 m/s error in the speed of sound of 1520 m/s causes a corresponding approximate 1% distortion in the apparent range, or 50 meters, i.e. $15/1520 \cdot 5000$.

4.5 Application of cMFP for tracking a moving source.

The advantage of cMFP over conventional MFP for locating a moving source along a long track is illustrated here. Fig. 10 displays the arbitrary path of a source moving along a parabolic trajectory (dashed lines). For the sake of simplicity, the Doppler effect is not accounted: this moving source scenario is simply simulated as 100 successive stationary sources located along the parabolic trajectory. For each positions, the SNR of the received signals at the vertical line array is constant and equal to 16 dB (Fig. 10.a) or 8 dB (Fig. 10.b). Conventional broadband coherent MFP is implemented by running 100 successive backpropagations per frequency over the search grid to estimate the source trajectory (see crosshair symbols). On the other hand, broadband coherent cMFP is implemented using $M = 2$ random backpropagations per frequency to estimate the same source trajectory (see cross symbols). The median value of the distance errors (computed from Eq. (31)) between the estimated and actual source trajectory is 1m when using both MFP and cMFP for a SNR of 16dB. A slightly higher error of 1.6m (resp. 1.1m) for the cMFP (resp. MFP) was found for a SNR of 8dB. Overall, Fig. 10 indicates that cMFP can potentially achieve comparable source tracking performance with a significantly reduced number of simulations.

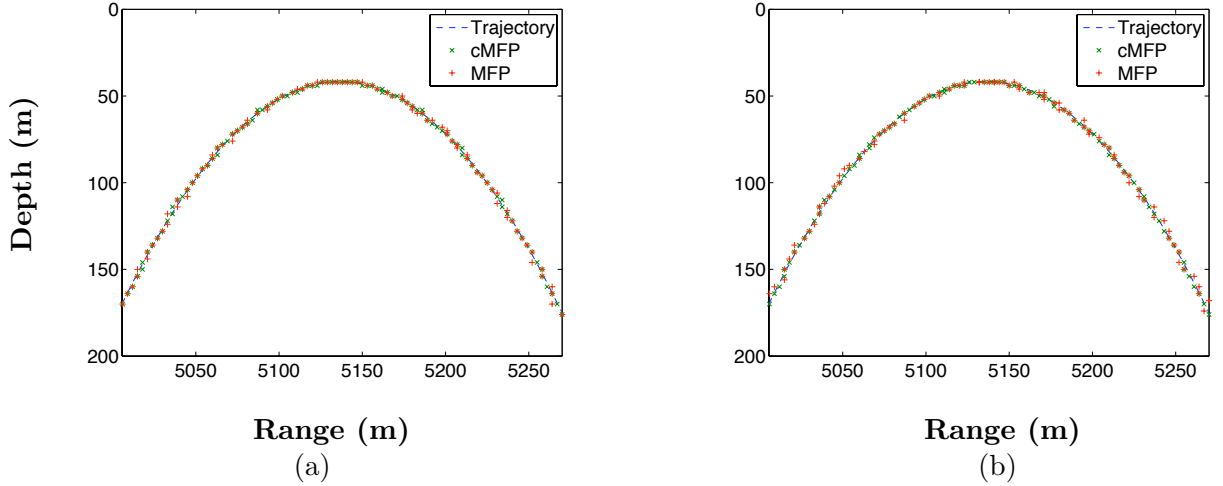


Figure 10: Tracking of a source moving along a parabolic source trajectory (dashed line) using either coherent broadband cMFP, implemented with $M = 2$ random backpropagations per frequency for the whole search grid, or using conventional broadband coherent MFP. For each of the 100 source positions, the SNR of the received signals at the vertical line array is constant and equal to (a) 16dB or (b) 8 dB.

5 Extension to adaptive MFP

Several variants of the MFP algorithm have been proposed in the existing literature [25, 21] to enhance the robustness and performance of the basic Bartlett formulation presented above (see Eq. (6)). This can be especially beneficial in the presence of added coherent noise to the received data vector Y_ω (see Eq. (3)). To do so, these higher resolution MFP algorithms are data adaptive, but typically have also a high resolution in their environmental knowledge requirements. A commonly used adaptive MFP formulation is the Minimum Variance Distortionless Response (MVDR) formulation. The MVDR formulation adaptively constructs a replica (or weighting) vector to yield a minimum mean square response to the recorded noise field along the receiver array while maintaining a constraint of unity processing gain for the incoming signal vector Y_ω [21, pg 540–552]:

$$|h_\omega^{MVDR}(\vec{r})|^2 = \left((G_\omega(\vec{r}))^H \mathbf{K}^{-1} G_\omega(\vec{r}) \right)^{-1}. \quad (32)$$

where \mathbf{K} is the $N \times N$ empirical correlation matrix from multiple realizations of the noisy received data vector Y_ω :

$$\mathbf{K} = \sum_{l=1}^L Y_{\omega,l} Y_{\omega,l}^H. \quad (33)$$

The physical interpretation and performance analysis of the MVDR formulation (see Eq. (32)) over the simple Bartlett formulation (see Eq. (6)) have been discussed extensively in the previous literature [25, 21] and thus will not be further repeated in this article.

The previous cMFP formulation can be readily extended to handle adaptive variants of the simple Bartlett MFP algorithm as discussed in Section III.C. For instance, using Eq. (32) and by direct analogy to Eq. (15), the magnitude square of the compressive MVDR

ambiguity surface is:

$$|\tilde{h}_\omega^{MVDR}(\vec{r})|^2 = \left((\Phi G_\omega(\vec{r}))^H (\Phi \mathbf{K} \Phi^H)^{-1} \Phi G_\omega(\vec{r}) \right)^{-1}. \quad (34)$$

So once we have computed the M tests measurements $\Phi G_\omega(\vec{r})$, they can be readily applied to either the compressive adaptive MFP formulation (see Eq. (34)) or the simple Bartlett formulation (see Eq. (15)) to locate the unknown source.

6 Conclusions

We have shown here how dimension-reducing random projections can greatly reduce the computational cost involved with source localization via matched-field processing. When compared to the location of the maximum of the ambiguity surface obtained from conventional MFP using N distributed receivers, the localization error achieved by cMFP scales down as square root of the number of random backpropagations M . The proposed cMFP formulation has also the added benefit to be able locate any source within the search grid area using only M random backpropagations, while conventional MFP would require at least N backpropagations to do the same. Thus cMFP provides an effective speedup factor of N/M per frequency, which can be significant when a large number of receivers N is available to locate a broadband source. Consequently this cMFP technique enables the ability to both broaden the search space and employ more sophisticated models of the Green's function, without introducing worries about sacrificing real-time performance.

This compressive approach is not limited to source localization, and could be extended to a more general type of machine learning problem when matches are evaluated via inner products (or equivalently via Euclidean norms). This type of approach has the potential to substantially decrease computational complexity in these cases, while admitting a negligibly small probability of error.

A Closest point on a line

For fixed vectors $U, V \in \mathbb{C}^n$, the following optimization program finds the closest point on the line spanned by v to u ,

$$\min_{\beta \in \mathbb{C}} \|U - \beta V\|^2.$$

The functional above attains its minimum value of

$$\|U\|^2 - \frac{|V^H U|^2}{\|V\|^2}$$

when

$$\beta = \frac{V^H U}{\|V\|^2}.$$

This fact can be verified by differentiating $\|U - \beta V\|^2$ with respect to the real and imaginary parts of β , and solving for value of β that makes them both equal to zero.

References

- [1] D. Achlioptas. Database-friendly random projections: Johnson-lindenstrauss with binary coins. *Journal of Computer and System Science*, 66:671–687, 2003.
- [2] A.B. Baggeroer, W.A. Kuperman, and P.N. Mikhalevsky. An overview of matched field methods in ocean acoustics. *IEEE Journal of Oceanic Engineering*, 18(4):401–424, 1993.
- [3] AB Baggeroer, WA Kuperman, and H. Schmidt. Matched field processing: Source localization in correlated noise as an optimum parameter estimation problem. *The Journal of the Acoustical Society of America*, 83:571, 1988.
- [4] R. G. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- [5] R. G. Baraniuk and M. B. Wakin. Random projections of smooth manifolds. *Foundations of Computational Mathematics*, 9(1):51–77, 2009.
- [6] E. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Math.*, 59(8):1207–1223, 2006.
- [7] E. Candès and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions on Information Theory*, 52(12):5406–5245, December 2006.
- [8] E. Candès and M. Wakin. An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25(2):21–30, March 2008.
- [9] E.J. Candès, J.K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207, 2006.
- [10] Lawrence Carin, Dehong Liu, and Bin Guo. In situ compressive sensing. *Inverse Problems*, 24(1):15–23, 2008.
- [11] V. Cevher, M.F. Duarte, and R.G. Baraniuk. Distributed target localization via spatial sparsity. *European Signal Processing Conference (EUSIPCO)*, 2008.
- [12] S. Chaillat and G. Biros. FalMS: A fast algorithm for the inverse medium problem with multiple frequencies and multiple sources for the scalar Helmholtz equation, 2010.
- [13] S. Dasgupta and A. Gupta. An elementary proof of a theorem of johnson and lindenstrauss. *Random Structures and Algorithms*, 22(1):60–65, 2003.
- [14] M.A. Davenport, P.T. Boufounos, M.B. Wakin, and R.G. Baraniuk. Signal processing with compressive measurements. *Selected Topics in Signal Processing, IEEE Journal of*, 4(2):445–460, 2010.
- [15] D.L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.

- [16] A. Fannjiang, P. Yan, and T. Strohmer. Compressed remote sensing of sparse objects. *SIAM Journal on Imaging Sciences*, 3(3):595–618, 2009.
- [17] F.K. Gruber, E.A. Marengo, and A.J. Devaney. Time-reversal imaging with multiple signal classification considering multiple scattering between the targets. *The Journal of the Acoustical Society of America*, 115:3042, 2004.
- [18] A.C. Gurbuz, J.H. McClellan, and V. Cevher. A compressive beamforming method. *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2008.
- [19] N. Halko, P.G. Martinsson, and J.A. Tropp. Finding structure with randomness: Stochastic algorithms for constructing approximate matrix decompositions. *SIAM Review, Survey and Review section*, 53(2):217–288, 2011.
- [20] D. Healy and D. J. Brady. Compression at the physical interface. *IEEE Signal Processing Magazine*, 25(2):67–71, March 2008.
- [21] Finn B. Jensen, William A. Kuperman, M. B. Porter, and Henrik Schmidt. *Computational Ocean Acoustics*. AIP Press, 2000.
- [22] W. B. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Conference in Modern Analysis and Probability*, pages 189–206, 1984.
- [23] EA Marengo, RD Hernandez, YR Citron, FK Gruber, M. Zambrano, and H. Lev-Ari. Compressive sensing for inverse scattering. *XXIX URSI General Assembly, Chicago, Illinois*, 2008.
- [24] J. Romberg. Imaging via compressive sampling. *IEEE Signal Processing Magazine*, pages 14–20, March 2008.
- [25] A. Tolstoy. Applications of matched-field processing to inverse problems in underwater acoustics. *Inverse Problems*, 16:1655, 2000.
- [26] M.B. Wakin. Manifold-based signal recovery and parameter estimation from compressive measurements. *Arxiv preprint arXiv:1002.1247*, 2010.